Logical Methods in Computer Science
Vol. 4 (3:1) 2008, pp. 1–21
www.lmcs-online.org

Submitted Sep. 24, 2007
Published Jul. 29, 2008

ENRICHED μ–CALCULI MODULE CHECKING *

ALESSANDRO FERRANTE [a], ANIELLO MURANO [b], AND MIMMO PARENTE [c]

[a,c] Università di Salerno, Via Ponte don Melillo, 84084 - Fisciano (SA), Italy
e-mail address: {ferrante,parente}@dia.unisa.it

[b] Università di Napoli "Federico II", Dipartimento di Scienze Fisiche, 80126 Napoli, Italy
e-mail address: murano@na.infn.it

Abstract. The model checking problem for open systems has been widely studied in the literature, for both finite–state (*module checking*) and infinite–state (*pushdown module checking*) systems, with respect to *CTL* and *CTL**. In this paper, we further investigate this problem with respect to the μ-calculus enriched with nominals and graded modalities (*hybrid graded μ-calculus*), in both the finite–state and infinite-state settings. Using an automata-theoretic approach, we show that *hybrid graded μ-calculus module checking* is solvable in exponential time, while *hybrid graded μ-calculus pushdown module checking* is solvable in double-exponential time. These results are also tight since they match the known lower bounds for *CTL*. We also investigate the module checking problem with respect to the hybrid graded μ-calculus enriched with inverse programs (*Fully enriched μ-calculus*): by showing a reduction from the tiling problem, we show its undecidability. We conclude with a short overview of the model checking problem for the Fully enriched μ-calculus and the fragments obtained by dropping at least one of the additional constructs.

1. Introduction

Model-checking is a formal method, applied in system design, to automatically verify the ongoing behavior of *reactive systems* ([CE81, QS81]). In this verification technique the behavior of a system, formally described by a mathematical model, is checked against a behavioral constraint, usually specified by a formula in an appropriate temporal logic (for a survey, see [CGP99]).

In the process of modeling a system, we distinguish between *closed* and *open* systems [HP85]. While the behavior of a closed system is completely determined by the state of the system, the behavior of an open system depends on the ongoing interaction with its environment [Hoa85]. In model checking open systems, introduced and called *module-checking* in [KVW01], one should check the system with respect to arbitrary environments and should take into account uncertainty regarding the environment. In such a framework, the open finite–state system is described by a labeled state–transition graph, called in fact

1998 ACM Subject Classification: F.1.1, F.1.2, F.3.1, D.2.4.

Key words and phrases: Finite state machine, tree automaton, push down automaton, interactive and reactive computation, logics of programs, modal logic, μ-calculus, formal verification, model checking.

* The paper is based on [FM07] and [FMP07].

LOGICAL METHODS IN COMPUTER SCIENCE
DOI:10.2168/LMCS-4 (3:1) 2008
© A. Ferrante, A. Murano, and M. Parente
Creative Commons

module, whose set of states is partitioned into *system states* (where the system makes a transition) and *environment states* (where the environment makes a transition). Given a module $\mathcal{M}$, describing the system to be verified, and a temporal logic formula φ, specifying the desired behavior of the system, module checking asks whether for all possible environments, $\mathcal{M}$ satisfies φ. Therefore, in module checking it is not sufficient to check whether the full computation tree obtained by unwinding $\mathcal{M}$ (that corresponds to the interaction of $\mathcal{M}$ with a maximal environment) satisfies φ, but it is also necessary to verify that all trees obtained from the full computation tree by pruning some subtrees rooted in nodes corresponding to choices disabled by the environment (those trees represent the interactions of $\mathcal{M}$ with all the possible environments), satisfy φ. We collect all such trees in a set named $exec(\mathcal{M})$. It is worth noticing that each tree in $exec(\mathcal{M})$ represents a "memoryful" behavior of the environment. Indeed, the unwinding of a module $\mathcal{M}$ induces duplication of nodes, which allow different pruning of subtrees. To see an example, consider a two-drink dispenser machine that serves, upon customer request, tea or coffee. The machine is an open system and an environment for the system is an infinite line of thirsty people. Since each person in the line can prefer both tea and coffee, or only tea, or only coffee, each person suggests a different disabling of the external choices. Accordingly, there are many different possible environments to consider. In [KV97, KVW01], it has been shown that while for linear–time logics model and module checking coincide, module checking for specification given in *CTL* and *CTL** is exponentially harder than model checking in the size of the formula and preserves the linearity in the size of the model. Indeed, *CTL* and *CTL** module checking is Exptime–complete and 2Exptime–complete, respectively.

In [BMP05, AMV07], the module checking technique has been extended to infinite-state systems by considering *open pushdown systems* (*OPD*, for short). These are pushdown systems augmented with finite information that allows us to partition the set of configurations into *system* and *environment* configurations. To see an example of an open pushdown system, consider the above two-drink dispenser machine, with the additional constraint that a coffee can be served only if the number of coffees served up to that time is smaller than that of teas served. Such a machine can be clearly modeled as an open pushdown system (the stack is used to guarantee the inequality between served coffees and teas). In [BMP05], it has been shown that pushdown module checking is 2Exptime–complete for *CTL* and 3Exptime–complete for *CTL**. Thus, for pushdown systems, and for specification given in *CTL* and *CTL**, module checking is exponentially harder than model checking with respect to the size of the formula, while it preserves the exponential complexity with respect to the size of the model [Wal96, Wal00].

Among the various formalisms used for specifying properties, a valid candidate is the *μ-calculus*, a very powerful propositional modal logic augmented with least and greatest fixpoint operators [Koz83] (for a recent survey, see also [BS06]). The *Fully enriched μ–calculus* [BP04] is the extension of the μ–calculus with *inverse programs*, *graded modalities*, and *nominals*. Intuitively, inverse programs allow us to travel backwards along accessibility relations [Var98], nominals are propositional variables interpreted as singleton sets [SV01], and graded modalities enable statements about the number of successors of a state [KSV02]. By dropping at least one of the additional constructs, we get a *fragment* of the Fully enriched μ-calculus. In particular, by inhibiting backward modalities we get the fragment we call *hybrid graded μ-calculus*. In [BP04], it has been shown that satisfiability is undecidable in the *Fully enriched μ–calculus*. On the other hand, it has been shown in [SV01, BLMV06] that satisfiability for each of its fragments is decidable and Exptime-complete (for more

details, see also [BLMV08]). The upper bound result is based on an automata–theoretic approach via *two-way graded alternating parity tree automata (*2GAPT*)*, along with the fact that each fragment of the Fully enriched μ-calculus enjoys the *quasi-forest model property*. Intuitively, $2GAPT$ generalize alternating automata on infinite trees as inverse programs and graded modalities enrich the standard μ–calculus: $2GAPT$ can move up to a node's predecessor and move down to *at least* n or *all but* n successors. Moreover, a quasi-forest is a forest where nodes can have roots as successors and having quasi-forest model property means that any satisfiable formula has a quasi-forest as model. Using $2GAPT$ and the quasi-forest model property, it has been shown in [SV01, BLMV06] that given a formula φ of a fragment of the Fully enriched μ-calculus, it is possible to construct a $2GAPT$ accepting all trees encodings[1] quasi-forests modeling φ. Then, the exponential-upper bound follows from the fact that the emptiness problem for $2GAPT$ is solvable in Ptime [KPV02].

In this paper, we further investigate the module checking problem and its infinite-state extension, with respect to the *hybrid graded μ-calculus*. To see an example of module checking a finite-state open system w.r.t. an hybrid graded μ-calculus specification, consider again the above two-drink dispenser machine with the following extra feature: whenever a customer can choose a drink, he can also call the customer service or the security service. Suppose also that by taking one of these two new choices, the drink-dispenser machine stops dispensing drinks, up to the moment the customer finishes operating with the service. Assume that, for the labeled state–transition graph modeling the system, we label by *choose* the choosing state and by the nominals o_c and o_s the states in which the interaction with the customer and the security services start, respectively. Moreover, suppose we want to check the following property: "whenever the customer comes at a choice, he can choose for both the customer and the security services". This property can be formalized by the hybrid graded μ–calculus formula $\nu x.((choose \rightarrow \langle 1, call\rangle\ (o_c \vee o_s)) \wedge [0,-]x)$, which reads "it is always true that whenever the drink–dispenser is in the *choose* state, there are at least 2 *call*–successors in which $(o_c \vee o_s)$ holds". Clearly, the considered open system does not satisfy this formula. Indeed, it is not satisfied by the particular behavior that chooses always the same service.

By exploiting an automata–theoretic approach via tree automata, we show that hybrid graded μ–calculus module checking is decidable and solvable in Exptime in the size of the formula and Ptime in the size of the system. Thus, as in general, we pay an exponential–time blowup with respect to the model checking problem (and only w.r.t. the size of the formula) for the module checking investigation. In particular, we reduce the addressed module checking problem to the emptiness problem for graded alternating parity tree automata ($GAPT$). In more details, given a model $\mathcal{M}$ and an hybrid graded μ-calculus formula φ, we first construct in polynomial time a Büchi tree automaton (NBT) $\mathcal{A}_{\mathcal{M}}$ accepting $exec(M)$. The construction of $\mathcal{A}_{\mathcal{M}}$ we propose here extends that used in [KVW01] by also taking into account that $\mathcal{M}$ must be unwound in a quasi-forest, rather than a tree, with both nodes and edges labeled. Thus, the set $exec(\mathcal{M})$ is a set of quasi-forests, and the automaton $\mathcal{A}_{\mathcal{M}}$ we construct will accept all trees encodings of all quasi-forests of $exec(\mathcal{M})$. From the formula side, accordingly to [BLMV06], we can construct in a polynomial time a $GAPT$ $\mathcal{A}_{\not\models\varphi}$ accepting all models that do not satisfy φ, with the intent to check that none of these models are in $exec(\mathcal{M})$. Thus, we check that $\mathcal{M}$ models φ for every possible choice of the

[1]Encoding is done by using a new root node that connects all roots of the quasi-forest and new atomic propositions which are used to encode programs and successor nodes corresponding to nominals.

environment by checking whether the $\mathcal{L}(\mathcal{A}_{\mathcal{M}}) \cap \mathcal{L}(\mathcal{A}_{\not\models\varphi})$ is empty. The results follow from the fact that an NBT is a particular case of $GAPT$, which are closed under intersection and have the emptiness problem solvable in Exptime [BLMV06]. We also show a lower bound matching the obtained upper bound by using a reduction from the module checking for CTL, known to be Exptime-hard.

By exploiting again an automata-theoretic approach, we show that hybrid graded μ-calculus pushdown module checking is decidable and solvable in 2Exptime in the size of the formula and Exptime in the size of the system. Thus, as in general, with respect to the finite–state model checking case we pay an exponential–time blowup in the size of both the system and the formula for the use of pushdown systems, and an another exponential–time blowup in the size of the formula for the module checking investigation. Our approach allow us do not take the trivial 2Exptime result on both the size of the system and the formula, which can be easily obtained by combining the algorithms existing in the literature along with that one we introduce in this paper for the finite–state case. We solve the hybrid graded μ-calculus pushdown module checking by using a reduction to the emptiness problem for nondeterministic pushdown parity tree automata (PD–NPT). The algorithm we propose extends that given for the finite-state case. In particular, given an OPD $\mathcal{S}$, a module $\mathcal{M}$ induced by the configurations of $\mathcal{S}$, and an hybrid graded μ-calculus formula φ, we first construct in polynomial time a pushdown Büchi tree automaton (PD–NBT) $\mathcal{A}_{\mathcal{M}}$, accepting $exec(\mathcal{M})$. From the formula side, accordingly to [BLMV06], we can construct in a polynomial time a $GAPT$ $\mathcal{A}_{\not\models\varphi}$ accepting all models that do not satisfy φ. Thus, we can check that $\mathcal{M}$ models φ for every possible choice of the environment by checking whether $\mathcal{L}(\mathcal{A}_{\mathcal{M}}) \cap \mathcal{L}(\mathcal{A}_{\not\models\varphi})$ is empty. By showing a non-trivial exponential reduction of $2GAPT$ into NPT, we show a 2Exptime upper bound for the addressed problem. Since the pushdown module checking problem for CTL is 2Exptime-hard, we get that the addressed problem is then 2Exptime-complete.

As regarding the Fully enriched μ-calculus, we also investigate the module checking problem in a "rewind" framework in the following sense. As far as backward modalities concern, everytime the system goes back to an environment's node, he is always able to redefine a new pruning choice. Given a module $\mathcal{M}$ and a Fully enriched μ-calculus formula φ, we solve the rewind module checking problem by checking that all trees in $exec(\mathcal{M})$, always taking the same choice in duplicate environment nodes, satisfy φ. By showing a reduction from the tiling problem [Ber66], we show that the addressed problem is undecidable.

We conclude the paper with short considerations on the model checking on all of the fragments of the Fully enriched μ-calculus. In particular we show the problem to be Exptime-complete for a pushdown system which is allowed to push one symbol per time onto the stack, with respect to any fragment not including the graded modality: for the fragments with the graded modality, we show a 2Exptime upper bound.

The rest of the paper is organized as follows. In Section 2, we give all the necessary preliminaries, Section 3 contains the definition of module checking w.r.t. hybrid graded μ-calculus, and Section 4 contains definitions and known results about $2GAPT$ and PD–NPT. In Sections 5 and 6, we give our main results on module checking for the hybrid graded μ-calculus. In Section 7, we show the undecidability result for the Fully enriched module checking and conclude in Section 8 with some complexity considerations on model checking with all the fragments of the Fully enriched μ-calculus.

2. Preliminaries

In this section, we recall definitions of labeled forests and hybrid graded μ–calculus. We refer to [BLMV06] for more technical definitions and motivating examples.

2.1. **Labeled Forests.** For a finite set X, we denote the *size* of X by $|X|$, the set of words over X by X^*, the empty word by ε, and with X^+ we denote $X^* \setminus \{\varepsilon\}$. Given a word w in X^* and a symbol a of X, we use $w \cdot a$ to denote the word wa. Let $\mathbb{N}$ be the set of positive integers. For $n \in \mathbb{N}$, let $\mathbb{N}$ denote the set $\{1, 2, \ldots, n\}$. A *forest* is a set $F \subseteq \mathbb{N}^+$ such that if $x \cdot c \in F$, where $x \in \mathbb{N}^+$ and $c \in \mathbb{N}$, then also $x \in F$. The elements of F are called *nodes*, and words consisting of a single natural number are *roots* of F. For each root $r \in F$, the set $T = \{r \cdot x \mid x \in \mathbb{N}^* \text{ and } r \cdot x \in F\}$ is a *tree* of F (the tree *rooted at* r). For $x \in F$, the nodes $x \cdot c \in F$ where $c \in \mathbb{N}$ are the *successors* of x, denoted $sc(x)$, and x is their *predecessor*. The number of successors of a node x is called the *degree* of x ($deg(x)$). The degree h of a forest F is the maximum of the degrees of all nodes in F and the number of roots. A forest with degree h is an *h-ary* forest. A full h-ary forest is a forest having h roots and all nodes with degree h.

Let $F \subseteq \mathbb{N}^+$ be a forest, x a node in F, and $c \in \mathbb{N}$. As a convention, we take $x \cdot \varepsilon = \varepsilon \cdot x = x$, $(x \cdot c) \cdot -1 = x$, and $c \cdot -1$ as undefined. We call x a *leaf* if it has no successors. A *path* π in F is a word $\pi = x_1 x_2 \ldots$ of F such that x_1 is a root of F and for every $x_i \in \pi$, either x_i is a leaf (i.e., π ends in x_i) or x_i is a predecessor of x_{i+1}. Given two alphabets Σ_1 and Σ_2, a (Σ_1, Σ_2)–labeled forest is a triple $\langle F, V, E\rangle$, where F is a forest, $V : F \to \Sigma_1$ maps each node of F to a letter in Σ_1, and $E : F \times F \to \Sigma_2$ is a partial function that maps each pair (x, y), with $y \in sc(x)$, to a letter in Σ_2. As a particular case, we consider a forest without labels on edges as a Σ_1–labeled forest $\langle F, V\rangle$, and a *tree* as a forest containing exactly one tree. A *quasi–forest* is a forest where each node may also have roots as successors. For a node x of a quasi–forest, we set $children(x)$ as $sc(x) \setminus \mathbb{N}$. All the other definitions regarding forests easily extend to quasi–forests. Notice that in a quasi–forest, since each node can have a root as successor, a root can also have several predecessors, while every other node has just one. Clearly, a quasi–forest can always be transformed into a forest by removing root successors.

2.2. **Hybrid Graded μ–Calculus.** Let AP, Var, $Prog$, and Nom be finite and pairwise disjoint sets of *atomic propositions*, *propositional variables*, *atomic programs* (which allow to travel the system along accessibility relations), and *nominals* (which are particular atomic propositions interpreted as singleton sets). The set of *hybrid graded μ–calculus* formulas is the smallest set such that

- **true** and **false** are formulas;
- p and $\neg p$, for $p \in AP$, are formulas;
- o and $\neg o$, for $o \in Nom$, are formulas;
- $x \in Var$ is a formula;
- if φ_1 and φ_2 are formulas, $\alpha \in Prog$, n is a non negative integer, and $y \in Var$, then the following are also formulas:

$$\varphi_1 \vee \varphi_2, \varphi_1 \wedge \varphi_2, \langle n, \alpha\rangle\varphi_1, [n, \alpha]\varphi_1, \ \mu y.\varphi_1(y), \text{ and } \nu y.\varphi_1(y).$$

Observe that we use positive normal form, i.e., negation is applied only to atomic propositions.

We call μ and ν *fixpoint operators*. A propositional variable y occurs *free* in a formula if it is not in the scope of a fixpoint operator. A *sentence* is a formula that contains no free variables. We refer often to the *graded modalities* $\langle n, \alpha\rangle\varphi_1$ and $[n, \alpha]\varphi_1$ as respectively *atleast formulas* and *allbut formulas* and assume that the integers in these operators are given in binary coding: the contribution of n to the length of the formulas $\langle n, \alpha\rangle\varphi$ and $[n, \alpha]\varphi$ is $\lceil \log n \rceil$ rather than n.

The semantics of the hybrid graded μ–calculus is defined with respect to a *Kripke structure*, i.e., a tuple $\mathcal{K} = \langle W, W_0, R, L\rangle$ where W is a non–empty set of *states*, $W_0 \subseteq W$ is the set of initial states, $R : Prog \to 2^{W\times W}$ is a function that assigns to each atomic program a transition relation over W, and $L : AP \cup Nom \to 2^W$ is a labeling function that assigns to each atomic proposition and nominal a set of states such that the sets assigned to nominals are singletons and subsets of W_0. If $(w, w') \in R(\alpha)$, we say that w' is an α–*successor* of w. Informally, an *atleast* formula $\langle n, \alpha\rangle\varphi$ holds at a state w of $\mathcal{K}$ if φ holds in at least $n + 1$ α–successors of w. Dually, the *allbut* formula $[n, \alpha]\varphi$ holds in a state w of $\mathcal{K}$ if φ holds in all but at most n α–successors of w. Note that $\neg\langle n, \alpha\rangle\varphi$ is equivalent to $[n, \alpha]\neg\varphi$, and the modalities $\langle\alpha\rangle\varphi$ and $[\alpha]\varphi$ of the standard μ–calculus can be expressed as $\langle 0, \alpha\rangle\varphi$ and $[0, \alpha]\varphi$, respectively.

To formalize semantics, we introduce valuations. Given a Kripke structure $\mathcal{K} = \langle W$, W_0, R, $L\rangle$ and a set $\{y_1, \ldots, y_n\}$ of variables in Var, a *valuation* $\mathcal{V} : \{y_1, \ldots, y_n\} \to 2^W$ is an assignment of subsets of W to the variables $y_1, \ldots, y_n$. For a valuation $\mathcal{V}$, a variable y, and a set $W' \subseteq W$, we denote by $\mathcal{V}[y \leftarrow W']$ the valuation obtained from $\mathcal{V}$ by assigning W' to y. A formula φ with free variables among $y_1, \ldots, y_n$ is interpreted over $\mathcal{K}$ as a mapping $\varphi^{\mathcal{K}}$ from valuations to 2^W, i.e., $\varphi^{\mathcal{K}}(\mathcal{V})$ denotes the set of points that satisfy φ under valuation $\mathcal{V}$. The mapping $\varphi^{\mathcal{K}}$ is defined inductively as follows:

- $\mathbf{true}^{\mathcal{K}}(\mathcal{V}) = W$ and $\mathbf{false}^{\mathcal{K}}(\mathcal{V}) = \emptyset$;
- for $p \in AP \cup Nom$, we have $p^{\mathcal{K}}(\mathcal{V}) = L(p)$ and $(\neg p)^{\mathcal{K}}(\mathcal{V}) = W \setminus L(p)$;
- for $y \in Var$, we have $y^{\mathcal{K}}(\mathcal{V}) = \mathcal{V}(y)$;
- $(\varphi_1 \wedge \varphi_2)^{\mathcal{K}}(\mathcal{V}) = \varphi_1^{\mathcal{K}}(\mathcal{V}) \cap \varphi_2^{\mathcal{K}}(\mathcal{V})$ and $(\varphi_1 \vee \varphi_2)^{\mathcal{K}}(\mathcal{V}) = \varphi_1^{\mathcal{K}}(\mathcal{V}) \cup \varphi_2^{\mathcal{K}}(\mathcal{V})$;
- $(\langle n, \alpha\rangle\varphi)^{\mathcal{K}}(\mathcal{V}) = \{w : |\{w' \in W : (w, w') \in R(\alpha) \text{ and } w' \in \varphi^{\mathcal{K}}(\mathcal{V})\}| \geq n + 1\}$;
- $([n, \alpha]\varphi)^{\mathcal{K}}(\mathcal{V}) = \{w : |\{w' \in W : (w, w') \in R(\alpha) \text{ and } w' \notin \varphi^{\mathcal{K}}(\mathcal{V})\}| \leq n\}$;
- $(\mu y.\varphi(y))^k(\mathcal{V}) = \bigcap\{W' \subseteq W : \varphi^{\mathcal{K}}([y \leftarrow W']) \subseteq W'\}$;
- $(\nu y.\varphi(y))^k(\mathcal{V}) = \bigcup\{W' \subseteq W : W' \subseteq \varphi^{\mathcal{K}}([y \leftarrow W'])\}$.

For a state w of a Kripke structure $\mathcal{K}$, we say that $\mathcal{K}$ *satisfies* φ at w if $w \in \varphi^{\mathcal{K}}$. In what follows, a formula φ *counts* up to b if the maximal integer in *atleast* and *allbut* formulas used in φ is $b - 1$.

3. Hybrid graded μ-calculus module Checking

In this paper we consider open systems, i.e., systems that interact with their environment and whose behavior depends on this interaction. The (global) behavior of such a system is described by a *module* $\mathcal{M} = \langle W_s, W_e, W_0, R, L\rangle$, which is a Kripke structure where the set of states $W = W_s \cup W_e$ is partitioned in *system states* W_s and *environment states* W_e.

Given a module $\mathcal{M}$, we assume that its states are ordered and the number of successors of each state w is finite. For each $w \in W$, we denote by $succ(w)$ the ordered tuple (possibly empty) of w's α-successors, for all $\alpha \in Prog$. When $\mathcal{M}$ is in a system state w_s, then all

states in $succ(w_s)$ are possible next states. On the other hand, when $\mathcal{M}$ is in an environment state w_e, the possible next states (that are in $succ(w_e)$) depend on the current environment. Since the behavior of the environment is not predictable, we have to consider all the possible sub–tuples of $succ(w_e)$. The only constraint, since we consider environments that cannot block the system, is that not all the transitions from w_e are disabled.

The set of all (maximal) computations of $\mathcal{M}$, starting from W_0, is described by a $(W, Prog)$–labeled quasi–forest $\langle F_\mathcal{M}, V_\mathcal{M}, E_\mathcal{M}\rangle$, called *computation quasi–forest*, which is obtained by unwinding $\mathcal{M}$ in the usual way. The problem of deciding, for a given branching–time formula φ over $AP \cup Nom$, whether $\langle F_\mathcal{M}, L \circ V_\mathcal{M}, E_\mathcal{M}\rangle$ satisfies φ at a root node, denoted $\mathcal{M} \models \varphi$, is the usual *model–checking problem* [CE81, QS81]. On the other hand, for an open system $\mathcal{M}$, the quasi–forest $\langle F_\mathcal{M}, V_\mathcal{M}, E_\mathcal{M}\rangle$ corresponds to a very specific environment, i.e., a maximal environment that never restricts the set of its next states. Therefore, when we examine a branching–time formula φ w.r.t. $\mathcal{M}$, the formula φ should hold not only in $\langle F_\mathcal{M}, V_\mathcal{M}, E_\mathcal{M}\rangle$, but in all quasi-forests obtained by pruning from $\langle F_\mathcal{M}, V_\mathcal{M}, E_\mathcal{M}\rangle$ subtrees rooted at children of environment nodes, as well as inhibiting some of their jumps to roots (that is, successor nodes labeled with nominals), if there are any. The set of these quasi–forests, which collects all possible behaviors of the environment, is denoted by $exec(\mathcal{M})$ and is formally defined as follows. A quasi–forest $\langle F, V, E\rangle \in exec(\mathcal{M})$ iff

- for each $w_i \in W_0$, we have $V(i) = w_i$;
- for each $x \in F$, with $V(x) = w$, $succ(w) = \langle w_1, \ldots, w_n, w_{n+1}, \ldots, w_{n+m}\rangle$, and $succ(w) \cap W_0 = \langle w_{n+1}, \ldots, w_{n+m}\rangle$, there exists $S = \langle w'_1, \ldots, w'_p, w'_{p+1}, \ldots, w'_{p+q}\rangle$ sub-tuple of $succ(w)$ such that $p + q \geq 1$ and the following hold:
  - $S = succ(w)$ if $w \in W_s$;
  - $children(x) = \{x \cdot 1, \ldots, x \cdot p\}$ and, for $1 \leq j \leq p$, we have $V(x \cdot j) = w'_j$ and $E(x, x \cdot j) = \alpha$ if $(w, w'_j) \in R(\alpha)$;
  - for $1 \leq j \leq q$, let $x_j \in \mathbb{N}$ such that $V(x_j) = w'_{p+j}$, then $E(x, x_j) = \alpha$ if $(w, w'_{p+j}) \in R(\alpha)$.

In the following, we consider quasi–forests in $exec(\mathcal{M})$ as labeled with $(2^{AP\cup Nom}, Prog)$, i.e., taking the label of a node x as $L(V(x))$. For a module $\mathcal{M}$ and a formula φ of the hybrid graded μ–calculus, we say that $\mathcal{M}$ *reactively* satisfies φ, denoted $\mathcal{M} \models_r \varphi$ (where "r" stands for *reactively*), if all quasi-forests in $exec(\mathcal{M})$ satisfy φ. The problem of deciding whether $\mathcal{M} \models_r \varphi$ is called *hybrid graded μ–calculus module checking*.

3.1. **Open Pushdown Systems (*OPD*).** An *OPD* over AP, Nom and $Prog$ is a tuple $\mathcal{S} = \langle Q, \Gamma, \flat, C_0, \Delta, \rho_1, \rho_2, Env\rangle$, where Q is a finite set of (control) *states*, Γ is a finite *stack alphabet*, $\flat \notin \Gamma$ is the *stack bottom symbol*. We set $\Gamma_\flat = \Gamma \cup \{\flat\}$, $Conf = Q \times (\Gamma^* \cdot \flat)$ to be the set of *(pushdown) configurations*, and for each configuration $(q, A \cdot \gamma)$, we set $top((q, A \cdot \gamma)) = (q, A)$ to be a *top configuration*. The function $\Delta : Prog \to 2^{(Q\times\Gamma_\flat)\times(Q\times\Gamma^*_\flat)}$ is a finite set of transition rules such that $\flat$ is always present at the bottom of the stack and nowhere else (thus whenever $\flat$ is read, it is pushed back). Note that we make this assumption also about the various pushdown automata we use later. The set $C_0 \subseteq Conf$ is a finite set of *initial configurations*, $\rho_1 : AP \to 2^{Q\times\Gamma_\flat}$ and $\rho_2 : Nom \to C_0$ are labeling functions associating respectively to each atomic proposition p a set of top configurations in which p holds and to each nominal exactly one initial configuration. Finally, $Env \subseteq Q \times \Gamma_\flat$ specifies the set of *environment configurations*. The size $|\mathcal{S}|$ of $\mathcal{S}$ is $|Q| + |\Delta| + |\Gamma|$.

The *OPD* moves in accordance with the transition relation Δ. Thus, $((q, A), (q', \gamma)) \in \Delta(\alpha)$ implies that if the *OPD* is in state q and the top of the stack is A, it can move along with an α*–transition* to state q', and substitute γ for A. Also note that the possible operations of the system, the labeling functions, and the designation of configurations as environment configurations, are all dependent only on the current control state and the top of the stack.

An *OPD* $\mathcal{S}$ induces a module $\mathcal{M}_{\mathcal{S}} = \langle W_s, W_e, W_0, R, L\rangle$, where:

- $W_s \cup W_e = Conf$, i.e. the set of pushdown configurations, and $W_0 = C_0$;
- $W_e = \{c \in Conf \mid top(c) \in Env\}$.
- $((q, A \cdot \gamma), (q', \gamma' \cdot \gamma)) \in R(\alpha)$ *iff* there is $((q, A), (q', \gamma')) \in \Delta(\alpha)$;
- $L(p) = \{c \in Conf \mid top(c) \in \rho_1(p)\}$ for $p \in AP$; $L(o) = \rho_2(o)$ for $o \in Nom$.

The *hybrid graded (μ-calculus) pushdown module checking* problem is to decide, for a given *OPD* $\mathcal{S}$ and an enriched μ–calculus formula φ, whether $\mathcal{M}_{\mathcal{S}} \models_r \varphi$.

4. Tree Automata

4.1. **Two-way Graded Alternating Parity Tree Automata ($2GAPT$).** These automata have been introduced and deeply investigated in [BLMV06]. In this section we just recall the main definitions and results and refer to the literature for more details. Intuitively, $2GAPT$ are an extension of nondeterministic tree automata in such a way that a $2GAPT$ can send several copies of itself to the same successor (*alternating*), send copies of itself to the predecessor (*two-way*), specify a number n of successors to which copies of itself are sent (*graded*), and accept trees along with a *parity acceptance condition*. To give a more formal definition, let us recall some technicalities from [BLMV06].

For a given set Y, let $B^+(Y)$ be the set of positive Boolean formulas over Y (i.e., Boolean formulas built from elements in Y using $\wedge$ and $\vee$), where we also allow the formulas **true** and **false** and $\wedge$ has precedence over $\vee$. For a set $X \subseteq Y$ and a formula $\theta \in B^+(Y)$, we say that X satisfies θ iff assigning **true** to elements in X and assigning **false** to elements in $Y \setminus X$ makes θ **true**. For $b > 0$, let $\langle [b]\rangle = \{\langle 0\rangle, \langle 1\rangle, \ldots, \langle b\rangle\}$, $[[b]] = \{[0], [1], \ldots, [b]\}$, and $D_b = \langle [b]\rangle \cup [[b]] \cup \{-1, \varepsilon\}$. Intuitively, D_b collects all possible directions in which the automaton can proceed.

Formally, a $2GAPT$ on Σ-labeled trees is a tuple $\mathcal{A} = \langle \Sigma, b, Q, \delta, q_0, \mathcal{F}\rangle$, where Σ is the input alphabet, $b > 0$ is a counting bound, Q is a finite set of states, $\delta : Q \times \Sigma \to B^+(D_b \times Q)$ is a transition function, $q_0 \in Q$ is an initial state, and $\mathcal{F}$ is a parity acceptance condition (see below). Intuitively, an atom $(\langle n\rangle, q)$ (resp. $([n], q)$) means that $\mathcal{A}$ sends copies in state q to $n + 1$ (resp. all but n) different successors of the current node, (ε, q) means that $\mathcal{A}$ sends a copy (in state q) to the current node, and $(-1, q)$ means that $\mathcal{A}$ sends a copy to the predecessor of the current node. A *run* of $\mathcal{A}$ on an input Σ-labeled tree $\langle T, V\rangle$ is a tree $\langle T_r, r\rangle$ in which each node is labeled by an element of $T \times Q$. Intuitively, a node in T_r labeled by (x, q) describes a copy of the automaton in state q that reads the node x of T. Runs start in the initial state and satisfy the transition relation. Thus, a run $\langle T_r, r\rangle$ with root z has to satisfy the following: (i) $r(z) = (1, q_0)$ for the root 1 of T and (ii) for all $y \in T_r$ with $r(y) = (x, q)$ and $\delta(q, V(x)) = \theta$, there is a (possibly empty) set $S \subseteq D_b \times Q$, such that S satisfies θ, and for all $(d, s) \in S$, the following hold:

- If $d \in \{-1, \varepsilon\}$, then $x \cdot d$ is defined, and there is $j \in \mathbb{N}$ such that $y \cdot j \in T_r$ and $r(y \cdot j) = (x \cdot d, s)$;

- If $d = \langle n \rangle$, there are at least $n+1$ distinct indexes $i_1, \ldots, i_{n+1}$ such that for all $1 \leq j \leq n+1$, there is $j' \in \mathbb{N}$ such that $y \cdot j' \in T_r$, $x \cdot i_j \in T$, and $r(y \cdot j') = (x \cdot i_j, s)$.
- If $d = [n]$, there are at least $deg(x) - n$ distinct indexes $i_1, \ldots, i_{deg(x)-n}$ such that for all $1 \leq j \leq deg(x) - n$, there is $j' \in \mathbb{N}$ such that $y \cdot j' \in T_r$, $x \cdot i_j \in T$, and $r(y \cdot j') = (x \cdot i_j, s)$.

Note that if $\theta = \mathbf{true}$, then y does not need to have successors. This is the reason why T_r may have leaves. Also, since there exists no set S as required for $\theta = \mathbf{false}$, we cannot have a run that takes a transition with $\theta = \mathbf{false}$.

A run $\langle T_r, r \rangle$ is *accepting* if all its infinite paths satisfy the acceptance condition. In the parity acceptance condition, $\mathcal{F}$ is a set $\{F_1, \ldots, F_k\}$ such that $F_1 \subseteq \ldots \subseteq F_k = Q$ and k is called the *index* of the automaton. An infinite path π on T_r *satisfies* $\mathcal{F}$ if there is an even i such that π contains infinitely many states from F_i and finitely many states from F_{i-1}. An automaton *accepts* a tree iff there exists an accepting run of the automaton on the tree. We denote by $\mathcal{L}(\mathcal{A})$ the set of all Σ-labeled trees that $\mathcal{A}$ accepts. The *emptiness* problem for an automaton $\mathcal{P}$ is to decide whether $\mathcal{L}(\mathcal{P}) = \emptyset$.

A $2GAPT$ is a $GAPT$ (i.e., "*one–way*") if $\delta : Q \times \Sigma \to B^+(D_b \setminus \{-1\} \times Q)$ and a $2APT$ (i.e., "*non-graded*") if $\delta : Q \times \Sigma \to B^+(\{-1, \varepsilon, 1, \ldots, h\} \times Q)$. As a particular case of $2APT$, we also consider nondeterministic parity tree automata (NPT) [KVW00]. Formally, an NPT on Σ-labeled trees is a tuple $\mathcal{A} = \langle \Sigma, D, Q, \delta, q_0, \mathcal{F} \rangle$, where Σ, Q, q_0, and $\mathcal{F}$ are as in $2APT$, D is a finite set of *branching degree* and $\delta : Q \times \Sigma \times D \to 2^{Q^*}$ is a transition function satisfying $\delta(q, \sigma, d) \subseteq Q^d$, for each $q \in Q$, $\sigma \in \Sigma$, and $d \in D$. Finally, we also consider *Büchi acceptance condition* $\mathcal{F} \subseteq Q$, which simply is a special parity condition $\{\emptyset, \mathcal{F}, Q\}$. Thus, we use in the following the acronym NBT to denote nondeterministic Büchi tree automata on Σ-labeled trees.

The following results on $2GAPT$ will be useful in the rest of the paper.

Theorem 1. [BLMV06] *The emptiness problem for a* GAPT $\mathcal{A} = \langle \Sigma, b, Q, \delta, q_0, \mathcal{F} \rangle$ *can be solved in time linear in the size of* Σ *and* b*, and exponential in the index of the automaton and number of states.*

Lemma 1. [BLMV06] *Given two* GAPT $\mathcal{A}_1$ *and* $\mathcal{A}_2$*, there exists a* GAPT $\mathcal{A}$ *such that* $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{A}_1) \cap \mathcal{L}(\mathcal{A}_2)$ *and whose size is linear in the size of* $\mathcal{A}_1$ *and* $\mathcal{A}_2$*.*

We now recall a result on $GAPT$ and hybrid graded μ-calculus formulas.

Lemma 2 ([BLMV06])**.** *Given an hybrid graded* μ*-calculus sentence* φ *with* ℓ atleast *subsentences and counting up to* b*, it is possible to construct a* GAPT *with* $\mathcal{O}(|\varphi|^2)$ *states, index* $|\varphi|$*, and counting bound* b *that accepts exactly each tree that encodes a quasi-forest model of* φ*.*

4.2. **Nondeterministic Pushdown Parity Tree Automata (*PD–NPT*).** A *PD–NPT* (without ε-transitions), on Σ-labeled full h-ary trees, is a tuple $\mathcal{P} = \langle \Sigma, \Gamma, \natural, Q, q_0, \gamma_0, \rho, \mathcal{F} \rangle$, where Σ is a finite input alphabet, Γ, $\natural$, $\Gamma_\natural$, and Q are as in OPD, (q_0, γ_0) is the initial configuration, $\rho : Q \times \Sigma \times \Gamma_\natural \to 2^{(Q \times \Gamma_\natural^*)^h}$ is a transition function, and $\mathcal{F}$ is a parity acceptance condition over Q. Intuitively, when $\mathcal{P}$ is in state q, reading an input node x labeled by $\sigma \in \Sigma$, and the stack contains a word $A \cdot \gamma \in \Gamma^* \cdot \natural$, then $\mathcal{P}$ chooses a tuple $\langle (q_1, \gamma_1), \ldots, (q_h, \gamma_h) \rangle \in \rho(q, \sigma, A)$ and splits in h copies such that for each $1 \leq i \leq h$, a copy in configuration $(q_i, \gamma_i \cdot \gamma)$ is sent to the node $x \cdot i$ in the input tree. A run of $\mathcal{P}$ on a Σ-labeled full h-ary tree $\langle T, V \rangle$ is a $(Q \times \Gamma^* \cdot \natural)$-labeled tree $\langle T, r \rangle$ such that

- $r(\varepsilon) = (q_0, \gamma_0)$, and
- for each $x \in T$ with $r(x) = (q, A \cdot \gamma)$, there is $\langle (q_1, \gamma_1), \ldots, (q_h, \gamma_h) \rangle \in \rho(q, V(x), A)$ such that, for all $1 \leq i \leq h$, we have $r(x \cdot i) = (q_i, \gamma_i \cdot \gamma)$.

The notion of accepting path is defined with respect to the control states that appear infinitely often in the path (thus without taking into account any stack content). Then, the notions given for $2GAPT$ regarding accepting runs, accepted trees, and accepted languages, along with the parity acceptance condition, easily extend to *PD–NPT*. In the following, we denote with *PD–NBT* a *PD–NPT* with a Büchi condition. We now recall two useful results on the introduced automata.

Proposition 4.1 ([KPV02])**.** *The emptiness problem for a* PD–NPT *on* Σ*-labeled full* h*-ary trees, having index* m*,* n *states, and transition function* ρ*, can be solved in time exponential in* $n \cdot m \cdot h \cdot |\rho|$*.*

Proposition 4.2 ([BMP05])**.** *Given a* PD–NBT $\mathcal{P} = \langle \Sigma, \Gamma, Q, q_0, \gamma_0, \rho, Q \rangle$ *on* Σ*-labeled full* h*-ary trees, and an* NPT $\mathcal{A} = \langle \Sigma, Q', q'_0, \delta, \mathcal{F}' \rangle$*, there is a* PD–NPT $\mathcal{P}'$ *on* Σ*-labeled full* h*-ary trees, such that* $\mathcal{L}(\mathcal{P}') = \mathcal{L}(\mathcal{P}) \cap \mathcal{L}(\mathcal{A})$*. Moreover,* $\mathcal{P}'$ *has* $|Q| \cdot |Q'|$ *states, the same index as* $\mathcal{A}$*, and the size of the transition relation is bounded by* $|\rho| \cdot |\delta| \cdot h$*.*

5. Deciding Hybrid Graded μ-calculus Module Checking

In this section, we solve the module checking problem for the hybrid graded μ–calculus. In particular, we show that this problem is decidable and Exptime–complete. For the upper bound, we give an algorithm based on an automata–theoretic approach, by extending an idea of [KVW01]. For the lower bound, we give a reduction from the module checking problem for *CTL*, known to be Exptime–hard. We start with the upper bound.

Let $\mathcal{M}$ be a module and φ an hybrid graded μ–calculus formula. We decide the module checking problem for $\mathcal{M}$ against φ by building a $GAPT$ $\mathcal{A}_{\mathcal{M} \times \not\models \varphi}$ as the intersection of two automata. Essentially, the first automaton, denoted by $\mathcal{A}_{\mathcal{M}}$, is a Büchi automaton that accepts trees encoding of labeled quasi–forests of $exec(\mathcal{M})$, and the second automaton is a $GAPT$ $\mathcal{A}_{\not\models \varphi}$ that accepts all trees encoding of labeled quasi–forests that do not satisfy φ (i.e, $\neg\varphi$ is satisfied at all initial nodes). Thus, $\mathcal{M} \models_r \varphi$ iff $\mathcal{L}(\mathcal{A}_{\mathcal{M} \times \not\models \varphi})$ is empty.

The construction of $\mathcal{A}_{\mathcal{M}}$ proposed here extends that given in [KVW01] for solving the module checking problem for finite–state open systems with respect to *CTL* and *CTL**. The extension concerns the handling of forest models instead of trees and formulas of the hybrid graded μ–calculus. Before starting, there are a few technical difficulties to be overcome. First, we notice that $exec(\mathcal{M})$ contains quasi–forests, with labels on both edges and nodes, while Büchi automata can only accept trees with labels on nodes. This problem is overcome by using the following three step transformation

(1) move the label of each edge to the target node of the edge (formally using a new propositional symbol p_α, for each atomic program α),
(2) substitute edges to roots with new propositional symbols $\uparrow_o^\alpha$ (which represents an α–labeled edge from the current node to the unique root node labeled by the nominal o), and
(3) add a new root, labeled with a new symbol *root*, and connect it with the old roots of the quasi–forest.

Let $AP' = AP \cup \{p_\alpha \mid \alpha \in Prog\} \cup \{\uparrow_o^\alpha \mid \alpha \in Prog \text{ and } o \in Nom\}$, we denote with $\langle T, V'\rangle$ the $(2^{AP' \cup Nom} \cup \{root\})$–labeled tree encoding of a quasi–forest $\langle F, V, E\rangle \in exec(\mathcal{M})$, obtained using the above transformation.

Another technical difficulty to handle is relate to the fact that quasi–forests of $exec(\mathcal{M})$ (and thus their encoding) may not share the same structure, since they are obtained by pruning some subtrees from the computation quasi–forest $\langle F_\mathcal{M}, V_\mathcal{M}, E_\mathcal{M}\rangle$ of $\mathcal{M}$. Let $\langle T_\mathcal{M}, V'_\mathcal{M}\rangle$ the *computation tree* of $\mathcal{M}$ obtained from $\langle F_\mathcal{M}, V_\mathcal{M}, E_\mathcal{M}\rangle$ using the above encoding. By extending an idea of [KVW01], we solve the technical problem by considering each tree $\langle T, V'\rangle$, encoding of a quasi–forest of $exec(\mathcal{M})$, as a $(2^{AP' \cup Nom} \cup \{root, \bot\})$–labeled tree $\langle T_\mathcal{M}, V''\rangle$ (where $\bot$ is a fresh proposition name not belonging to $AP \cup Nom \cup \{root\}$) such that for each node $x \in T_\mathcal{M}$, if $x \in T$ then $V''(x) = V'(x)$, otherwise $V''(x) = \{\bot\}$. Thus, we label each node pruned in the $\langle T_\mathcal{M}, V'_\mathcal{M}\rangle$ with $\{\bot\}$ and recursively, we label with $\{\bot\}$ its subtrees. In this way, all trees encoding quasi–forests of $exec(\mathcal{M})$ have the same structure of $\langle T_\mathcal{M}, V'_\mathcal{M}\rangle$, and they differ only in their labeling.

Accordingly, we can think of an environment as a strategy for placing $\{\bot\}$ in $\langle T_\mathcal{M}, V'_\mathcal{M}\rangle$, with the aim of preventing the system to satisfy a desired property while not considering the nodes labeled with $\bot$. Moreover, the environment can also disable jumps to roots. This is performed by removing from nodes corresponding with environment states some of $\uparrow_o^\alpha$ labels. Notice that since we consider environments that do not block the system, each node associated with an environment state has at least one successor not labeled by $\{\bot\}$, unless it has $\uparrow_o^\alpha$ in its label.

Let us denote by $\widehat{exec}(\mathcal{M})$ the set of all $(2^{AP' \cup Nom} \cup \{root, \bot\})$–labeled $\langle T_\mathcal{M}, V''\rangle$ trees obtained from $\langle F, V, E\rangle \in exec(\mathcal{M})$ in the above described manner. The required NBT $\mathcal{A}_\mathcal{M}$ must accept all and only the $(2^{AP' \cup Nom} \cup \{root, \bot\})$–labeled trees in $\widehat{exec}(\mathcal{M})$. The automaton $\mathcal{A}_\mathcal{M} = \langle \Sigma, D, Q, \delta, q_0, \mathcal{F}\rangle$ is defined for a module $\mathcal{M} = \langle W_s, W_e, W_0, R, L\rangle$ as follows:

- $\Sigma = 2^{AP' \cup Nom} \cup \{root, \bot\}$
- $D = \bigcup_{w \in W} |succ(w) \setminus W_0|$ (that is, D contains, for each state in W, the number of its successors, but its jumps to roots).
- $Q = (W \times \{\bot, \top, \vdash\}) \cup \{q_0\}$, with $q_0 \notin W$. Thus every state w of $\mathcal{M}$ induces three states $(w, \bot)$, $(w, \top)$, and $(w, \vdash)$ in $\mathcal{A}_\mathcal{M}$. Intuitively, when $\mathcal{A}_\mathcal{M}$ is in state $(w, \bot)$, it can read only $\bot$, in state $(w, \top)$, it can read only letters in $2^{AP^* \cup\ Nom}$, and in state $(w, \vdash)$, it can read both letters in $2^{AP^* \cup Nom}$ and $\bot$. In this last case, it is left to the environment to decide whether the transition to a state of the form $(w, \vdash)$ is enabled. The three types of states are used to ensure that the environment enables all transitions from enabled system states, enables at least one transition from each enabled environment state, and disables transitions from disabled states.
- The transition function $\delta : Q \times \Sigma \times D \to 2^{Q^*}$ is defined as follows. Let $x \in T$ be a node of the input tree.
  - if $root \in V(x)$ then (let $W_0 = \{w_1, \ldots, w_m\}$)
    $$\delta(q_0, root, m) = \{\langle (w_1, \top), \ldots, (w_m, \top)\rangle\},$$
    that is $\delta(q_0, root, m)$ contains exactly one m–tuple of all the roots of the forest. In this case, all transitions cannot be disabled;
  - if $root \notin V(x)$, let $V(x) = w$ and $succ(w) \setminus W_0 = \langle w_1, \ldots, w_n\rangle$ be the set of *non–roots successors* of w, then we have

* for $w \in W_e \cup W_s$ and $g \in \{\vdash, \bot\}$ we have
$$\delta((w,g), \bot, n) = \{\langle (w_1, \bot), \ldots, (w_n, \bot)\rangle\},$$
that is $\delta((w,g), \bot, n)$ contains exactly one n–tuple of all non–roots successors of w. In this case, all transitions to successors of w are recursively disabled;
* for $w \in W_s$ and $g \in \{\top, \vdash\}$ we have
$$\delta((w,g), L(w), n) = \{\langle (w_1, \top), \ldots, (w_n, \top)\rangle\},$$
that is, $\delta((w,g), L(w), n)$ contains exactly one n–tuple of all non–roots successors of w. In this case all transitions to successors of w are enabled;
* for $w \in W_e$ and $g \in \{\top, \vdash\}$ with $L(w) \cap \{\uparrow_o^\alpha |\ \alpha \in Prog \text{ and } o \in Nom\} = \emptyset$ (i.e., w has no jumps to roots or all of them have been disabled), we have
$$\begin{aligned}\delta((w,g), L(w), n) = \{\ & \langle (w_1, \top),\ (w_2, \vdash), \ldots, (w_n, \vdash)\rangle,\\ & \langle (w_1, \vdash),\ (w_2, \top), \ldots, (w_n, \vdash)\rangle,\\ & \qquad\qquad \vdots\\ & \langle (w_1, \vdash),\ (w_2, \vdash), \ldots, (w_n, \top)\rangle\},\end{aligned}$$
that is, $\delta((w,g), L(w), n)$ contains n different n–tuples of all non–roots successors of w. When $\mathcal{A}_\mathcal{M}$ proceeds according to the i–th tuple, the environment can disable all transitions to successors of w, except that to w_i;
* for $w \in W_e$ and $g \in \{\top, \vdash\}$ with $L(w) \cap \{\uparrow_o^\alpha |\ \alpha \in Prog \text{ and } o \in Nom\} \neq \emptyset$ (i.e., w has at least one jump to roots enabled), we have
$$\delta((w,g), L(w), n) = \{\langle (w_1, \vdash), \ldots, (w_n, \vdash)\rangle\},$$
that is $\delta((w,g), L(w), n)$ contains one n–tuple of non–roots successors of w, that can be successively disabled.

Notice that δ is not defined when n is different from the number of non–roots successors of w, and when the input does not meet the restriction imposed by the $\top$, $\vdash$, and $\bot$ annotations or by the labeling of w.

The automaton $\mathcal{A}_\mathcal{M}$ has $3 \cdot |W| + 1$ states, $2^{|AP| \cdot |R|} + 2$ symbols, and the size of the transition relation $|\delta|$ is bounded by $|R|(|W| \cdot 2^{|R|})$.

We recall that a node labeled by either $\{\bot\}$ or $\{root\}$ stands for a node that actually does not exist. Thus, we have to take this into account when we interpret formulas of the hybrid graded μ–calculus over trees $\langle T_\mathcal{M}, V'\rangle \in \widehat{exec}(\mathcal{M})$. In order to achieve this, as in [KVW01] we define a function f that transforms the input formula φ in a formula of the hybrid graded μ–calculus $\varphi' = \langle 0, \alpha\rangle f(\varphi)$ (where $\alpha \in Prog$ is an arbitrary atomic program), that restricts path quantification to only paths that never visit a state labeled with $\{\bot\}$. The function f we consider extends that given in [KVW01] and is inductively defined as follows:

- $f(\mathbf{true}) = \mathbf{true}$ and $f(\mathbf{false}) = \mathbf{false}$;
- $f(p) = p$ and $f(\neg p) = \neg p$ for all $p \in AP \cup Nom$;
- $f(x) = x$ for all $x \in Var$;
- $f(\varphi_1 \vee \varphi_2) = f(\varphi_1) \vee f(\varphi_2)$ and $f(\varphi_1 \wedge \varphi_2) = f(\varphi_1) \wedge f(\varphi_2)$ for all hybrid graded μ–calculus formulas φ_1 and φ_2;
- $f(\mu x.\varphi(x)) = \mu x.f(\varphi(x))$ and $f(\nu x.\varphi(x)) = \nu x.f(\varphi(x))$ for all $x \in Var$ and hybrid graded μ–calculus formulas φ;

- $f(\langle n,\alpha\rangle\varphi) = \langle n,\alpha\rangle(\neg\bot \wedge f(\varphi))$ for $n \in \mathbb{N}$ and for all atomic programs α and hybrid graded μ–calculus formulas φ;
- $f([n,\alpha]\varphi) = [n,\alpha](\neg\bot \wedge f(\varphi))$ for $n \in \mathbb{N}$ and for all atomic programs α and hybrid graded μ–calculus formulas φ.

By definition of f, it follows that for each formula φ and $\langle T,V\rangle \in \widehat{exec}(\mathcal{M})$, $\langle T,V\rangle$ satisfies $\varphi' = \langle 0,\alpha\rangle f(\varphi)$ *iff* the $2^{AP' \cup Nom}$–labeled forest, obtained from $\langle T,V\rangle$ removing the node labeled with $\{root\}$ and all nodes labeled by $\{\bot\}$, satisfies φ. Therefore, we solve the module checking problem of $\mathcal{M}$ against an hybrid graded μ–calculus formula φ by checking (for its negation) that in $\widehat{exec}(\mathcal{M}) = \mathcal{L}(\mathcal{A}_{\mathcal{M}})$ does not exist any tree $\langle T,V\rangle$ satisfying $\neg\varphi' = [0,\alpha]f(\neg\varphi)$ (note that $|f(\neg\varphi)| = \mathcal{O}(|\neg\varphi|)$). We reduce the latter to check the emptiness of a $GAPT$ $\mathcal{A}_{\mathcal{M}\times\not\models f(\varphi)}$ that is defined as the intersection of the NBT $\mathcal{A}_{\mathcal{M}}$ with a $GAPT$ $\mathcal{A}_{\not\models f(\varphi)}$ accepting exactly the $2^{AP' \cup Nom} \cup \{root, \bot\}$ trees encodings of quasi–forests not satisfying $f(\varphi)$. By Lemma 2, if φ is an hybrid graded μ–calculus formula, then $\mathcal{A}_{\not\models f(\varphi)}$ has $\mathcal{O}(|\varphi|^2)$ states, index $|\varphi|$, and counting bound b. Therefore, by Lemma 1, $\mathcal{A}_{\mathcal{M}\times\not\models f(\varphi)}$ has $\mathcal{O}(|W| + |\varphi|^2)$ states, index $|\varphi|$, and counting bound b. By recalling that the emptiness problem for a $GAPT$ can be decided in exponential-time (Theorem 1), we obtain that the module checking problem for hybrid graded μ–calculus formulas is solvable in exponential-time. To show a tight lower bound we recall that CTL module checking is EXPTIME–hard [KVW01] and every CTL formula can be linearly transformed in a modal μ–calculus formula [Jur98]. This leads to the module checking problem w.r.t. modal μ–calculus formulas to be EXPTIME–hard and thus to the following result.

Theorem 2. *The module checking problem with respect to hybrid graded μ–calculus formulas is* EXPTIME*–complete.*

6. Deciding Hybrid Graded μ-calculus PD-module Checking

In this section, we show that hybrid graded pushdown module checking is decidable and solvable in 2EXPTIME. Since CTL pushdown module checking is 2EXPTIME–hard, we get that the addressed problem is 2EXPTIME–complete. For the upper bound, the algorithm works as follows. Given an OPD $\mathcal{S}$ and the module $\mathcal{M}_{\mathcal{S}}$ induced by $\mathcal{S}$, by combining and extending the constructions given in [BMP05] and Section 5, we first build in polynomial–time a PD–NBT $\mathcal{A}_{\mathcal{S}}$ accepting each tree that encodes a quasi–forest belonging to $exec(\mathcal{M}_{\mathcal{S}})$. Then, given an hybrid graded μ–calculus formula φ, according to [BLMV06], we build in polynomial–time a $GAPT$ $\mathcal{A}_{\not\models\varphi}$ (Lemma 2) accepting all models that do not satisfy φ, with the intent of checking that none of these models are in $exec(\mathcal{M}_{\mathcal{S}})$. Then, accordingly to the basic idea of [KVW01], we check that $\mathcal{M}_{\mathcal{S}} \models_r \varphi$ by checking whether $\mathcal{L}(\mathcal{A}_{\mathcal{S}}) \cap \mathcal{L}(\mathcal{A}_{\not\models\varphi})$ is empty. Finally, we get the result by using an exponential–time reduction of the latter to the emptiness problem for PD–NPT, which from Proposition 4.1 can be solved in EXPTIME. As a key step of the above reduction, we use the exponential–time translation from $GAPT$ into NPT showed in Lemma 5.

Let us start dealing with $\mathcal{A}_{\mathcal{S}}$. Before building the automaton, there are some technical difficulties to overcome. First, notice that $\mathcal{A}_{\mathcal{S}}$ is a PD–NBT and it can only deal with trees having labels on nodes. Also, quasi–forests of $exec(\mathcal{M}_{\mathcal{S}})$ may not share the same structure, since they are obtained by pruning subtrees from the computation quasi–forest $\langle F_{\mathcal{M}_{\mathcal{S}}}, V_{\mathcal{M}_{\mathcal{S}}}, E_{\mathcal{M}_{\mathcal{S}}}\rangle$ of $\mathcal{M}_{\mathcal{S}}$. As in Section 5, we solve this problem by considering $2^{AP' \cup Nom} \cup$

$\{root, \bot\}$–labeled trees *encoding* of quasi–forests $\langle F, V, E\rangle \in exec(\mathcal{M}_\mathcal{S})$, where $AP' = AP \cup \{p_\alpha \mid \alpha \in Prog\} \cup \{\uparrow_o^\alpha \mid \alpha \in Prog \text{ and } o \in Nom\}$.

Another technical difficulty to handle with is related to the fact that quasi–forests of $exec(\mathcal{M}_\mathcal{S})$ (and thus their encodings) may not be full h–ary, since the nodes of the *OPD* from which $\mathcal{M}_\mathcal{S}$ is induced may have different degrees. Technically, we need this property since the emptiness problem for *PD–NPT* to which we reduce our problem has been solved in the literature only for *PD–NPT* working on full trees. Similarly as we did for pruned nodes, we transform each tree encoding of a quasi–forest of $exec(\mathcal{M}_\mathcal{S})$ into a full h–ary tree by adding missing nodes labeled with $\{\bot\}$. Therefore the proposition $\bot$ is used to denote both "disabled" states and "completion" states. In this way, all trees encodings of quasi–forests of $exec(\mathcal{M}_\mathcal{S})$ are all full h–ary trees, and they differ only in their labeling. Let us denote with $\widehat{exec}(\mathcal{M}_\mathcal{S})$ the set of all $2^{AP' \cup Nom} \cup \{root, \bot\}$–labeled full h–ary trees obtained from $\langle F_{\mathcal{M}_\mathcal{S}}, V_{\mathcal{M}_\mathcal{S}}, E_{\mathcal{M}_\mathcal{S}}\rangle$ using all the transformations described above.

In [BMP05] it has been shown how to build a *PD–NBT* accepting full h–ary trees embedded in an *OPD* corresponding to all behaviors of the environment. In particular, the *PD–NBT* constructed there already takes into account the above transformation regarding $\{\bot\}$–labeled nodes. By extending the construction proposed there in the same way the construction showed in Section 5 extends the classical construction of $\mathcal{A}_\mathcal{M}$ proposed in [KVW01], it is not hard to show that the following result holds.

Lemma 3. *Given an* OPD $\mathcal{S} = \langle Q, \Gamma, \flat, C_0, \Delta, \rho_1, \rho_2, Env\rangle$ *with branching degree* h*, we can build a* PD–NBT $\mathcal{A}_\mathcal{S} = \langle \Sigma, \Gamma, \flat, Q', q_0', \gamma_0, \delta, Q\rangle$*, which accepts exactly* $\widehat{exec}(\mathcal{M}_\mathcal{S})$*, such that* $\Sigma = 2^{AP' \cup Nom} \cup \{root, \bot\}$*,* $|Q'| = \mathcal{O}(|Q|^2 \cdot |\Gamma|)$*, and* $|\delta|$ *is polynomially bounded by* $h \cdot |\Delta|$*.*

Let us now go back to the hybrid graded μ–calculus formula φ. Using the function f introduced in Section 5 and Lemma 2, we get that given an hybrid graded μ–calculus formula φ, we can build in polynomial–time a *GAPT* $\mathcal{A}_{\not\models f(\varphi)}$ accepting all models of $\neg\varphi' = [0, \alpha] f(\neg\varphi)$ (as done in Section 5).

By using the classical Exptime transformation from *GAPT* to *GNPT* [KSV02] and a simple Exptime transformation from *GNPT* to *NPT*, we directly get a 3Exptime algorithm for the hybrid graded μ–calculus pushdown module checking. To obtain an exponential–time improvement, here we show a not trivial Exptime transformation from $2GAPT$ to *NPT*. The translation we propose uses the notions of *strategies*, *promises* and *annotations*, which we now recall.

Let $\mathcal{A} = \langle \Sigma, b, Q, \delta, q_0, \mathcal{F}\rangle$ be a $2GAPT$ with $\mathcal{F} = \langle F_1, \ldots, F_k\rangle$ and $\langle T, V\rangle$ be a Σ-labeled tree. Recall that $D_b = \langle [b]\rangle \cup [[b]] \cup \{-1, \varepsilon\}$ and $\delta : (Q \times \Sigma) \to B^+(D_b \times Q)$. For each control state $q \in Q$, let $index(q)$ be the minimal i such that $q \in F_i$. A *strategy tree* for $\mathcal{A}$ on $\langle T, V\rangle$ is a $2^{Q \times D_b \times Q}$-labeled tree $\langle T, \mathsf{str}\rangle$ such that, defined $head(w) = \{q : (q, d, q') \in w\}$ as the set of *sources* of w, it holds that (i) $q_0 \in head(\mathsf{str}(root(T)))$ and (ii) for each node $x \in T$ and state q, the set $\{(q, q') : (q, d, q') \in \mathsf{str}(x)\}$ satisfies $\delta(q, V(x))$.

A *promise tree* for $\mathcal{A}$ on $\langle T, V\rangle$ is a $2^{Q \times Q}$-labeled tree $\langle T, \mathsf{pro}\rangle$. We say that pro *fulfills* str for V if the states promised to be visited by pro satisfy the obligations induced by str as it runs on V. Formally, pro fulfills str for V if for every node $x \in T$, the following hold: "for every $(q, \langle n\rangle, q') \in \mathsf{str}(x)$ (resp. $(q, [n], q') \in \mathsf{str}(x)$), at least $n + 1$ (resp $deg(x) - n$) successors $x \cdot j$ of x have $(q, q') \in \mathsf{pro}(x \cdot j)$".

An *annotation tree* for $\mathcal{A}$ on $\langle T, \mathsf{str}\rangle$ and $\langle T, \mathsf{pro}\rangle$ is a $2^{Q \times \{1, \ldots, k\} \times Q}$-labeled tree $\langle T, \mathsf{ann}\rangle$ such that for each $x \in T$ and $(q, d_1, q_1) \in \mathsf{str}(x)$ the following hold:

- if $d_1 = \varepsilon$, then $(q, index(q_1), q_1) \in \mathsf{ann}(x)$;

- if $d_1 \in \{1,\dots,k\}$, then for all $d_2 \in \{1,\dots,k\}$ and $q_2 \in Q$ such that $(q_1,d_2,q_2) \in \mathsf{ann}(x)$, we have $(q,\min(d_1,d_2),q_2) \in \mathsf{ann}(x)$;
- if $d_1 = -1$ and $x = y \cdot i$, then for all $d_2,d_3 \in \{1,\dots,k\}$ and all $q_2,q_3 \in Q$ satisfing $(q_1,d_2,q_2) \in \mathsf{ann}(y)$ as well as $(q_2,d_3,q_3) \in \mathsf{str}(y)$ and $(q_2,q_3) \in \mathsf{pro}(x)$, we have that $(t,\min(index(q_1),d_2,index(q_3)),q_3) \in \mathsf{ann}(x)$;
- if $d_1 \in [[b]] \cup \langle [b] \rangle$, $y = x \cdot i$, and $(q,q_1) \in \mathsf{pro}(y)$, then for all $d_2,d_3 \in \{1,\dots,k\}$ and $q_2,q_3 \in Q$ such that $(q_1,d_2,q_2) \in \mathsf{ann}(y)$ and $(q_2,-1,q_3) \in \mathsf{str}(y)$, it holds that $(t,\min(index(q_1),d_2,index(q_3)),q_3) \in \mathsf{ann}(x)$.

A downward path induced by str, pro, and ann on $\langle T,V\rangle$ is a sequence $\langle x_0,q_0,t_0\rangle$, $\langle x_1,q_1,t_1\rangle$, ... such that $x_0 = root(T)$, q_0 is the initial state of $\mathcal{A}$ and, for each $i \geq 0$, it holds that $x_i \in T$, $q_i \in Q$, and $t_i = \langle q_i,d,q_{i+1}\rangle \in \mathsf{str}(x_i) \cup \mathsf{ann}(x_i)$ is such that either *(i)* $d \in \{1,\dots,k\}$ and $x_{i+1} = x_i$, or *(ii)* $d \in \langle [b] \rangle \cup [[b]]$ and there exists $c \in \{1,\dots,deg(x_i)\}$ such that $x_{i+1} = x_i \cdot c$ and $(q_i,q_{i+1}) \in \mathsf{pro}(x_{i+1})$. In the first case we set $index(t_i) = d$ and in the second case we set $index(t_i) = \min\{j \in \{1,\dots,k\} \mid q_{i+1} \in F_j\}$. Moreover, for a downward path π, we set $index(\pi)$ as the minimum index that appears infinitely often in π. Finally, we say that π is *accepting* if $index(\pi)$ is even.

The following lemma relates languages accepted by $2GAPT$ with strategies, promises, and annotations.

Lemma 4 ([BLMV06])**.** *Let $\mathcal{A}$ be a* 2GAPT*. A Σ-labeled tree $\langle T,V\rangle$ is accepted by $\mathcal{A}$ iff there exist a strategy tree $\langle T,\mathsf{str}\rangle$, a promise tree $\langle T,\mathsf{pro}\rangle$ for $\mathcal{A}$ on $\langle T,V\rangle$ such that* pro *fulfills* str *for V, and an annotation tree $\langle T,\mathsf{ann}\rangle$ for $\mathcal{A}$ on $\langle T,V\rangle$, $\langle T,\mathsf{str}\rangle$ and $\langle T,\mathsf{pro}\rangle$ such that every downward path induced by* str*,* pro*, and* ann *on $\langle T,V\rangle$ is accepting.*

Given an alphabet Σ for the input tree of a $2GAPT$ with transition function δ, let D_b^δ be the subset containing only the elements of D_b appearing in δ. Then we denote by Σ' the extended alphabet for the combined trees, i.e., $\Sigma' = \Sigma \times 2^{Q \times D_b^\delta \times Q} \times 2^{Q \times Q} \times 2^{Q \times \{1,\dots k\} \times Q}$.

Lemma 5. *Let $\mathcal{A}$ be a* 2GAPT *running on Σ–labeled trees with n states, index k and counting bound b that accepts h-ary trees. It is possible to construct in exponential-time an* NPT *$\mathcal{A}'$ running on Σ'–labeled h-ary trees that accepts a tree iff $\mathcal{A}$ accepts its projection on Σ.*

Proof. Let $\mathcal{A} = \langle \Sigma,b,Q,q_0,\delta,\mathcal{F}\rangle$ with $\mathcal{F} = \langle F_1,\dots,F_k\rangle$. By Lemma 4, we construct $\mathcal{A}'$ as the intersection of two NBT $\mathcal{A}'$, $\mathcal{A}''$, and an NPT $\mathcal{A}'''$. In particular, all these automata have size exponential in the size of $\mathcal{A}$. Moreover, since each NBT uses as accepting all its states, it is easy to intersect in polynomial-time all of them by using a classical automata product. These automata are defined as follows. Given a Σ'-labeled tree $T' = \langle T,(V,\mathsf{str},\mathsf{pro},\mathsf{ann})\rangle$,

(1) $\mathcal{A}'$ accepts T' iff str is a strategy for $\mathcal{A}$ on $\langle T,V\rangle$ and pro fulfills str for V,
(2) $\mathcal{A}''$ accepts T' iff ann is an annotation for $\mathcal{A}$ on $\langle T,V\rangle$, $\langle T,\mathsf{str}\rangle$ and $\langle T,\mathsf{pro}\rangle$, and
(3) $\mathcal{A}'''$ accepts T' iff every downward path induced by str, pro, and ann on $\langle T,V\rangle$ is accepting.

The automaton $\mathcal{A}' = \langle \Sigma',D',Q',q_0',\delta',\mathcal{F}'\rangle$ works as follows: on reading a node x labeled $(\sigma,\eta,\rho,\omega)$, then it locally checks whether η satisfies the definition of strategy for $\mathcal{A}$ on $\langle T,V\rangle$. In particular, when $\mathcal{A}'$ is in its initial state, we check that η contains a transition starting from the initial state of $\mathcal{A}$. Moreover, the automaton $\mathcal{A}'$ sends to each child $x \cdot i$ the pairs of states that have to be contained in $\mathsf{pro}(x \cdot i)$, in order to verify that pro fulfills str. To obtain this, we set $Q' = 2^{Q \times Q} \cup \{q_0'\}$, $D' = \{1,\dots,h\}$ and $\mathcal{F}' = \{\emptyset,Q'\}$. To define δ', we

first give the following definition. For each node $x \in T$ labeled $(\sigma, \eta, \rho, \omega)$, we set

$$\begin{array}{rcl} S(\eta) & = & \{\langle S_1, \ldots, S_{deg(x)}\rangle \in (2^{Q\times Q})^{deg(x)} \text{ such that} \\ & & [\text{for each } (q, \langle m\rangle, p) \in \eta \text{ there is } P \subseteq \{1, \ldots deg(x)\} \text{ with } |P| = m+1 \\ & & \text{such that for all } i \in P,\ (q,p) \in S_i] \text{ and} \\ & & [\text{for each } (q, [m], p) \in \eta \text{ there is} P \subseteq \{1, \ldots deg(x)\} \text{ with } |P| = deg(x) - m \\ & & \text{such that for all } i \in P,\ (q,p) \in S_i]\} \end{array}$$

to be the set of all tuples with size $deg(x)$, each *fulfilling* all graded modalities in $\mathsf{str}(x)$. Notice that $|S(\eta)| \leq 2^{hn^2}$. Then we have

$$\delta'(q, (\sigma, \eta, \rho, \omega), deg(x)) = \begin{cases} S(\eta) & \text{if } \forall\ p \in head(\eta),\ \{(d,p') \mid (p,d,p') \in \eta\} \text{ satisfies } \delta(p,\sigma) \\ & \text{and } [(q = q_0^1 \text{ and } q_0 \in head(\eta)) \text{ or } (q \neq q_0^1 \text{ and } q \subseteq \rho)] \\ \mathbf{false} & \text{otherwise.} \end{cases}$$

Hence, in $\mathcal{A}'$ we have $|Q'| = 2^{n^2}$, $|\delta'| \leq 2^{n^2(k+1)}$, and index 2.

$\mathcal{A}'' = \langle \Sigma', D'', Q'', q_0'', \delta'', \mathcal{F}''\rangle$ works in a similar way to $\mathcal{A}'$. That is, for each node x, it first locally checks whether the constraints of the annotations are verified; then it sends to the children of x the strategy and annotation associated with x, in order to successively verify whether the promises associated with the children nodes are consistent with the annotation of x. Therefore, in A'' we have $Q'' = 2^{Q\times D_b^\delta \times Q} \times 2^{Q\times\{1,\ldots,k\}\times Q}$, $q_0'' = (\emptyset, \emptyset)$, $\mathcal{F}'' = \{\emptyset, Q''\}$, $D'' = \{1, \ldots, h\}$, and for a state $(\eta_{prev}, \omega_{prev})$ and a letter $(\sigma, \eta, \rho, \omega)$ we have

$$\delta''((\eta_{prev}, \omega_{prev}), (\sigma, \eta, \rho, \omega), deg(x)) = \begin{cases} \langle (\eta, \omega), \ldots, (\eta, \omega)\rangle & \text{if the local conditions for the} \\ & \text{annotations are verified} \\ \mathbf{false} & \text{otherwise.} \end{cases}$$

Hence, in $\mathcal{A}''$ we have $|Q''| \leq 2^{n^2(|\delta|+k)}$, $|\delta''| \leq h \cdot 2^{n^2(|\delta|+k)}$, and index 2.

Finally, to define $\mathcal{A}'''$ we start by constructing a $2APT$ $\mathcal{B}$ whose size is polynomial in the size of $\mathcal{A}$ and accepts $\langle T, (V, \mathsf{str}, \mathsf{pro}, \mathsf{ann})\rangle$ iff there is a non accepting downward path (w.r.t. $\mathcal{A}$) induced by str, pro, and ann on $\langle T, V\rangle$. The automaton $\mathcal{B} = \langle \Sigma', Q^B, q_0^B, \delta^B, \mathcal{F}^B\rangle$ (which in particular does not need direction -1) essentially chooses, in each state, the downward path to walk on, and uses an integer to store the index of the state. We use a special state $\sharp$ not belonging to Q to indicate that $\mathcal{B}$ proceeds in accordance with an annotation instead of a strategy. Therefore, $Q^B = ((Q \cup \{\sharp\}) \times \{1, \ldots, k\} \times Q) \cup \{q_0^B\}$.

To define the transition function on a node x, let us introduce a function f that for each $q \in Q$, strategy $\eta \in 2^{Q\times D_b^\delta \times Q}$, and annotation $\omega \in 2^{Q\times\{1,\ldots,k\}\times Q}$ gives a formula satisfied along downward paths consistent with η and ω, starting from a node reachable in $\mathcal{A}$ with the state q. That is, in each node x, the function f either proceeds according to the annotation ω or the strategy η (note that f does not check that the downward path is consistent with any promise). Formally, f is defined as follows, where $index(p)$ is the minimum i such that $p \in F_i$:

$$f(q,\eta,\omega) = \bigvee_{\substack{(q,d,p)\in\omega \\ d\in\{1,\ldots,k\}}} \langle \varepsilon, (\sharp, d, p)\rangle \quad \vee \bigvee_{\substack{(q,d,p)\in\eta \\ d\in\langle[b]\rangle\cup[[b]]}} \bigvee_{c\in\{1,\ldots,deg(x)\}} \langle c, (q, index(p), p)\rangle$$

Then, we have $\delta^B(q_0^B, (\sigma, \eta, \rho, \omega)) = f(q_0, \eta, \omega)$ and

$$\delta^B((q,d,p), (\sigma,\eta,\rho,\omega)) \quad = \quad \begin{cases} \mathbf{false} & \text{if } q \neq \sharp \text{ and } (q,p) \notin \rho \\ f(p,\eta,\omega) & \text{otherwise.} \end{cases}.$$

A downward path π is non accepting for $\mathcal{A}$ if the minimum index that appears infinitely often in π is odd. Therefore, $\mathcal{F}^B = \langle F_1^B, \ldots, F_{k+1}^B, Q^B \rangle$ where $F_1^B = \emptyset$ and, for all $i \in \{2, \ldots, k+1\}$, we have $F_i^B = \{(q, d, p) \in Q^B \mid d = i - 1\}$. Thus, $|Q^B| = kn(n+1) + 1$, $|\delta^B| = k \cdot |\delta| \cdot |Q^B|$, and the index is $k + 2$. Then, since $\mathcal{B}$ is alternating, we can easily complement it in polynomial-time into a $2APT$ $\overline{\mathcal{B}}$ that accepts a tree iff all downward paths induced by str, pro, and ann on $\langle T, V \rangle$ are accepting. Finally, following [Var98] we construct in exponential-time the desired automaton $\mathcal{A}'''$. □

By applying the transformation given by Lemma 5 to the automaton $\mathcal{A}_{\neg\varphi'}$ defined above, we obtain in exponential time in the size of φ, an NPT that accepts all the trees encoding of quasi–forests that do not satisfy φ. From Proposition 4.2, then we can build a PD–NPT $\mathcal{A}_{\mathcal{S}\times\not\models\varphi}$ with size polynomial in the size of $\mathcal{S}$ and exponential in the size of φ such that $\mathcal{L}(\mathcal{A}_{\mathcal{S}\times\not\models\varphi}) = \mathcal{L}(\mathcal{A}_{\mathcal{S}}) \cap \mathcal{L}(\mathcal{A}_{\neg\varphi'})$. Hence, from Proposition 1 we obtain that hybrid graded μ–calculus pushdown module checking can be solved in EXPTIME in the size of $\mathcal{S}$ and in 2EXPTIME in the size of φ. Finally, from the fact that CTL pushdown module checking is known to be 2EXPTIME–hard with respect to the size of φ and EXPTIME–hard with respect to the size of $\mathcal{S}$ [BMP05], we obtain the following theorem.

Theorem 3. *The hybrid graded μ–calculus pushdown module checking problem is* 2EXPTIME*–complete with respect to the size of the formula and* EXPTIME*–complete with respect to the size of the system.*

7. Fully Enriched μ-calculus Module Checking

In this section, we consider a memoryless restriction of the module checking problem and investigate it with respect to formulas of the Fully enriched μ-calculus. Given a formula φ, a *memoryless module checking* problem checks whether all trees in $exec(\mathcal{M})$, always taking the same choice in duplicate environment nodes, satisfy φ. In this section, we show that the (memoryless) module checking problem for Fully enriched μ-calculus is undecidable.

Fully enriched μ–calculus is the extension of hybrid graded μ–calculus with *inverse programs*. Essentially, inverse programs allow us to specify properties about predecessors of a state. Given an atomic program $a \in Prog$, we denote its inverse program with a^- and the syntax of the fully enriched μ–calculus is simply obtained from the one we introduced for hybrid graded μ–calculus, by allowing both atomic and inverse programs in the graded modalities. Similarly, the semantics of fully enriched μ–calculus is given, identically to the one for hybrid graded μ–calculus, with respect to a Kripke structure $\mathcal{K} = \langle W, W_0, R, L \rangle$ in which, to deal with inverse programs, we define, for all $a \in Prog$, $R(a^-) = \{(v, w) \in W \times W$ such that $(w, v) \in R(a)\}$.

Let us note that, since the fully enriched μ–calculus does not enjoy the forest model property [BP04], we cannot unwind a Kripke structure in a forest. However, it is always possible to unwind it in an equivalent acyclic graph that we call *computation graph*. In order to take into account all the possible behaviors of the environment, we consider all the possible subgraphs of the computation graph obtained disabling some transitions from environment nodes but one. We denote with $graphs(M)$ the set of this graphs. Given a Fully enriched μ–calculus formula φ, we have that $M \models_r \varphi$ iff $K \models \varphi$ for all $K \in graphs(M)$.

To show the undecidability of the addressed problem, we need some further definitions. An (infinite) *grid* is a tuple $G = \langle \mathbb{N}^2, h, v \rangle$ such that h and v are defined as $h(\langle x, y \rangle) =$

$\langle x+1, y\rangle$ and $v(\langle x, y\rangle) = \langle x, y+1\rangle$. Given a finite set of *types* T, we will call *tile* on T a function $\hat{\rho} : \mathbb{N}^2 \to T$ that associates a type from T to each vertex of an infinite grid G, and we call *tiled infinite grid* the tuple $\langle G, T, \hat{\rho}\rangle$. A *grid model* is an infinite Kripke structure $K = \langle W, \{w_0\}, R, L\rangle$, on the set of atomic programs $Prog = \{l^-, v\}$, such that K can be mapped on a grid in such a way that w_0 corresponds to the vertex $\langle 0, 0\rangle$, $R(v)$ corresponds to v and $R(l^-)$ corresponds to h. We say that a grid model K "corresponds" to a tiled infinite grid $\langle G, T, \hat{\rho}\rangle$ if every state of K is labeled with only one atomic proposition (and zero or more nominals) and there exists a bijective function $\rho : T \to AP$ such that, if $w_{x,y}$ is the state of K corresponding with the node $\langle x, y\rangle$ of G, then $\rho(\hat{\rho}(\langle x, y\rangle)) \in L(w_{x,y})$.

Theorem 4. *The module checking problem for fully enriched μ–calculus is undecidable.*

Proof. To show the result, we use a reduction from the tiling problem (also known as *domino* problem), known to be undecidable [Ber66]. The tiling problem is defined as follows.

Let T be a finite set of types, and $H, V \subseteq T^2$ be two relations describing the types that cannot be vertically and horizontally adjacent in an infinite grid. The tiling problem is to decide whether there exists a tiled infinite grid $\langle G, T, \hat{\rho}\rangle$ such that $\hat{\rho}$ preserves the relations H and V. We call such a tile function a *legal tile* for G on T.

In [BP04], Bonatti and Peron showed undecidability for the satisfiability problem for fully enriched μ–calculus by also using a reduction from the tiling problem. Hence, given a set of types T and relations H and V, they build a (alternation free) fully enriched μ–calculus formula φ such that φ is satisfiable iff the tiling problem has a solution in a tiled infinite grid, with a legal tile ρ on T (with respect to H and V). In particular, the formula they build can be only satisfiable on a grid model K corresponding to a tiled infinite grid with a legal tile ρ on T. In the reduction we propose here, we use the formula φ used in [BP04]. It remains to define the module.

Let $\{G_1, G_2, \ldots\}$ be the set of all the infinite tiled grids on T (i.e., $G_i = \langle G, T, \hat{\rho}_i\rangle$), we build a module M such that $graphs(M)$ contains, for each $i \geq 1$, a grid models corresponding to G_i. Therefore, we can decide the tiling problem by checking whether $M \models_r \neg\varphi$. Indeed, if $M \models_r \neg\varphi$, then all grid models corresponding to G_i do not satisfy φ and, therefore, there is no solution for the tiling problem. On the other side, if $M \not\models_r \neg\varphi$, then there exists a model for φ; since φ can be satisfied only on a grid model corresponding to a tiled infinite grid with a legal tile on T with respect to H and V, we have that the tiling problem has a solution.

Formally, let $T = \{t_1, \ldots, t_m\}$ be the set of types, the module $M = \langle W_s, W_e, W_0, R, L\rangle$ with respect to atomic programs $Prog = \{l^-, v\}$, atomic propositions $AP = T$, and nominals $Nom = \{o_1, \ldots, o_m\}$, is defined as follows:

- $W_s = \emptyset$, $W_e = \{x_1, \ldots, x_m, y_1, \ldots, y_m\}$ and $W_0 = \{x_1, \ldots, x_m\}$;
- for all $i \in \{1, \ldots m\}$, $L(t_i) = \{x_i, y_i\}$ and $L(o_i) = \{x_i\}$;
- $R(v) = \{\langle x_i, x_j\rangle | i, j \in \{1, \ldots, m\}\} \cup \{\langle y_i, y_j\rangle | i, j \in \{1, \ldots, m\}\}$;
- $R(l^-) = \{\langle x_i, y_j\rangle | i, j \in \{1, \ldots, m\}\} \cup \{\langle y_i, x_j\rangle | i, j \in \{1, \ldots, m\}\}$

Notice that we duplicate the set of nodes labeled with tiles since we cannot have pairs of nodes in M labeled with more than one atomic program (in our case, with both v and l^-). Moreover the choice of labeling nodes x_i with nominals is arbitrary. Finally, from the fact that the module contains only environment nodes, it immediately follows that, for each i, the grid model corresponding to the infinite tiled grid G_i is contained in $graphs(M)$. □

μ-calculus extensions	Pushdown Model Checking	Single–Push Model Checking	Finite–State Model Checking
propositional	EXPTIME–Complete [Wal96]	EXPTIME–Complete	UP $\cap$ CO–UP [Wil01]
hybrid	EXPTIME–Complete	EXPTIME–Complete	UP $\cap$ CO–UP
graded	2EXPTIME	2EXPTIME	EXPTIME
full	?	EXPTIME–Complete	UP $\cap$ CO–UP
hybrid graded	2EXPTIME	2EXPTIME	EXPTIME
full hybrid	?	EXPTIME–Complete	UP $\cap$ CO–UP
full graded	?	2EXPTIME	EXPTIME
Fully enriched	?	2EXPTIME	EXPTIME

Figure 1: Results on Model Checking Problem.

8. Notes on Fully Enriched μ-calculus Model Checking

In this section, for the sake of completeness, we investigate the model checking problems for Fully enriched μ-calculus and its fragments, for both pushdown and finite-states systems.

In particular, we first consider the model checking problem for formulas of the μ–calculus enriched with nominal (*hybrid μ–calculus*) or graded modalities (*graded μ–calculus*) or both, for pushdown systems (PDMC, for short) and finite states systems (FSMC, for short), i.e. Kripke structures. In particular, we show that for graded μ–calculus, PDMC is solvable in 2EXPTIME and FSMC is solvable in EXPTIME. Moreover we show that for hybrid μ–calculus PDMC is EXPTIME–complete and FSMC is in UP $\cap$ CO–UP, thus matching the known results for (propositional) μ–calculus model checking (see [Wal96] for PDMC and [Wil01] for FSMC), and that, for hybrid graded μ–calculus, PDMC is solvable in 2EXPTIME and FSMC is solvable in EXPTIME.

By considering also μ-calculus enriched (among the others) with inverse programs, we also consider PDMC w.r.t. a reduced pushdown system that, in each transition, can increase the size of the stack by at most one (*single–push system*). To this aim, we define a single–push system with three stack operations: for $A \in \Gamma$, $sub(A)$ changes the top of the stack into A, $push(A)$ pushes the symbol A on the top of the stack, and $pop()$ pops the top symbol of the stack. Formally, a single–push system $\mathcal{S}$ is a pushdown system in which the transition function is $\Delta : Prog \to 2^{(Q\times\Gamma_\flat)\times(Q\times\{sub,push,pop\}\times\Gamma)}$. For consistency reasons, we assume that if the top of the stack is $\flat$ then $sub(A) = push(A)$ and $pop()$ has no effect.

We call the model checking problem for single–push systems *single–push model checking* (SPMC, for short). In this case, we show that for full hybrid μ–calculus (μ-calculus enriched with inverse programs and nominals), SPMC is EXPTIME–complete and FSMC is in UP $\cap$ CO–UP, and that for Fully enriched and full graded μ–calculus (μ-calculus enriched with inverse programs and graded modalities), SPMC is in 2EXPTIME and FSMC is in EXPTIME. In Figure 8 we report known and new results on model checking problems for the Fully enriched μ–calculus and its fragments.

To prove our results, we simply rule out inverse programs and nominals from the input formula. In particular, we first observe that, from a model checking point of view, checking a formula with inverse programs on a graph (finite or infinite) is equivalent to check the formula in "forward" on the graph enriched with opposite edges. That is, we consider inverse programs in the formula as special atomic programs to be checked on the opposite edges we have added in the graph. Note that this observation does not apply to PDMC.

Indeed, to transform previous configurations to inverse next configurations, we need to limit the power of a PDMC to be single push. Thus, we obtain the following result.

Lemma 6. *Let X_μ be an enrichment of the μ-calculus with inverse programs. Then a* SPMC *(resp.,* FSMC*) w.r.t. X_μ can be reduced in linear time to* SPMC *(resp.,* FSMC*) w.r.t. X_μ without inverse programs.*

Proof. Here we only show the proof for FSMC since the one for SPMC is similar. Let $\mathcal{K} = \langle W, W_0, R, L\rangle$ be a model that uses atomic programs from $Prog$, and let φ be a formula of X_μ. Then, we define a new model $\mathcal{K}'$ and a new formula φ' as follows: $\mathcal{K}' = \langle W, W_0, R', L\rangle$ uses atomic programs from the set $Prog' = Prog \cup \{\hat{a} \text{ s.t. } a \in Prog\}$ (it doesn't use inverse programs) and has the transition relation defined as $R'(a) = R(a)$ and $R'(\hat{a}) = R(a^-)$ for all $a \in Prog$. On the other side, φ' is a formula of X_μ without inverse programs equal to φ except for the fact that a^- is changed into $\hat{a}$ for all $a \in Prog$. Thus it can be easily seen that $\mathcal{K} \models \varphi$ iff $\mathcal{K}' \models \varphi'$ and this completes the proof of this lemma. □

Furthermore, from the model checking point of view, one can consider each nominal in the input formula as a particular atomic proposition. Thus we obtain the following result.

Lemma 7. *Let X_μ be the μ-calculus enriched with nominals and possibly with graded modalities. Then* PDMC*,* SPMC *and* FSMC *w.r.t. X_μ can be respectively reduced in linear time to* PDMC*,* SPMC *and* FSMC *w.r.t. X_μ without nominals.*

Proof. In this case too, we show the proof only for FSMC. Let $\mathcal{K} = \langle W, W_0, R, L\rangle$ be a model that uses atomic propositions from AP and nominals from Nom, and let φ be a formula of X_μ. Then, we consider the new model $\mathcal{K}' = \langle W, W_0, R, L\rangle$ that uses atomic propositions from the set $AP' = AP \cup Nom$ ($\mathcal{K}'$ does not use nominals); moreover, let φ' be the formula φ interpreted as a formula of X_μ without nominals on the set of atomic propositions AP'. Then, it is easy to see that $\mathcal{K} \models \varphi$ iff $\mathcal{K}' \models \varphi'$. □

From Lemmas 6 and 7 and the fact that for propositional μ–calculus PDMC is Exptime–Complete [Wal96] and FSMC is in UP ∩ co–UP [Wil01], we directly have that hybrid μ–calculus PDMC is Exptime–Complete, (full) hybrid μ–calculus SPMC is solvable in Exptime and (full) hybrid μ–calculus FSMC is in UP ∩ co–UP. Now, in [Wal96] it has been showed that μ-calculus PDMC is Exptime–hard. The proof used there can be easily adapted to handle single–push systems without incurring in any complexity blowup. Thus, we obtain the following result.

Theorem 5. *Hybrid μ–calculus* PDMC *is* Exptime*–Complete, (full) hybrid μ–calculus* SPMC *is* Exptime*–Complete and (full) hybrid μ–calculus* FSMC *is in* UP ∩ co–UP.

Finally, from Lemmas 6 and 7 we have that hybrid graded μ–calculus PDMC can be reduced in linear time to graded μ–calculus PDMC, Fully enriched μ–calculus SPMC can be reduced in linear time to graded μ–calculus SPMC (note that SPMC is a special case of PDMC) and Fully enriched μ-calculus FSMC can be reduced in linear time to graded μ-calculus FSMC. Since model checking is a special case of module checking, from Theorems 2 and 3 we have the following result.

Theorem 6. PDMC *is solvable in* 2Exptime *for (hybrid) graded μ–calculus,* SPMC *is solvable in* 2Exptime *for Fully enriched μ-calculus and* FSMC *is solvable in* Exptime *for Fully enriched μ-calculus.*

References

[AMV07] A. Aminof, A. Murano, and M.Y. Vardi, *Pushdown module checking with imperfect information*, CONCUR '07, LNCS, vol. 4703, Springer-Verlag, 2007, pp. 461–476.

[Ber66] R. Berger, *The undecidability of the domino problem*, Mem. AMS **66** (1966), 1–72.

[BLMV06] P.A. Bonatti, C. Lutz, A. Murano, and M.Y. Vardi, *The complexity of enriched μ-calculi*, ICALP '06, LNCS, vol. 4052, 2006, pp. 540–551.

[BLMV08] ———, *The complexity of enriched μ-calculi*, To appear in Logical Methods in Computer Science (2008), 1–27, http://www.na.infn.it/∼murano/pubblicazioni/journal-version-enriched.pdf.

[BMP05] L. Bozzelli, A. Murano, and A. Peron, *Pushdown module checking*, LPAR, 2005, pp. 504–518.

[BP04] P.A. Bonatti and A. Peron, *On the undecidability of logics with converse, nominals, recursion and counting*, Artificial Intelligence **158 : 1** (2004), 75–96.

[BS06] J. Bradfield and C. Stirling, *Modal μ-calculi*, Handbook of Modal Logic (Blackburn, Wolter, and van Benthem, eds.), Elsevier, 2006, pp. 722–756.

[CE81] E.M. Clarke and E.A. Emerson, *Design and synthesis of synchronization skeletons using branching time temporal logic*, Proc. of Work. on Logic of Programs, LNCS, vol. 131, 1981, pp. 52–71.

[CGP99] E.M. Clarke, O. Grumberg, and D.A. Peled, *Model checking*, MIT Press, Cambridge, MA, USA, 1999.

[FM07] A. Ferrante and A. Murano, *Enriched μ–calculus module checking*, FOSSACS'07, LNCS, vol. 4423, 2007, p. 183197.

[FMP07] A. Ferrante, A. Murano, and M. Parente, *Enriched μ–calculus pushdown module checking*, LPAR'07, LNAI, vol. 4790, 2007, pp. 438–453.

[Hoa85] C.A.R. Hoare, *Communicating sequential processes*, 1985.

[HP85] D. Harel and A. Pnueli, *On the development of reactive systems*, Logics and Models of Concurrent Systems, NATO Advanced Summer Institutes, vol. F–13, Springer–Verlag, 1985, pp. 477–498.

[Jur98] Marcin Jurdziński, *Deciding the winner in parity games is in $UP \cap co\text{-}UP$*, Information Processing Letters **68** (1998), no. 3, 119–124.

[Koz83] D. Kozen, *Results on the propositional mu–calculus.*, Theoretical Computer Science **27** (1983), 333–354.

[KPV02] O. Kupferman, N. Piterman, and M.Y. Vardi, *Pushdown specifications*, LPAR, 2002, pp. 262–277.

[KSV02] O. Kupferman, U. Sattler, and M.Y. Vardi, *The complexity of the graded μ–calculus*, CADE '02, LNAI, vol. 2392, 2002, pp. 423–437.

[KV97] O. Kupferman and M.Y. Vardi, *Module checking revisited*, CAV '96, LNCS, vol. 1254, Springer-Verlag, 1997, pp. 36–47.

[KVW00] O. Kupferman, M.Y. Vardi, and P. Wolper, *An automata–theoretic approach to branching–time model checking*, Journal of ACM **47** (2000), no. 2, 312–360.

[KVW01] ———, *Module checking*, Information & Computation **164** (2001), 322–344.

[QS81] J.P. Queille and J. Sifakis, *Specification and verification of concurrent systems in cesar*, 5^{th} Symp. on Programming, LNCS, vol. 137, 1981, pp. 337–351.

[SV01] U. Sattler and M.Y. Vardi, *The hybrid mu–calculus*, IJCAR '01, LNAI, vol. 2083, 2001, pp. 76–91.

[Var98] M.Y. Vardi, *Reasoning about the past with two–way automata*, ICALP '98, LNCS, vol. 1443, 1998, pp. 628–641.

[Wal96] I. Walukiewicz, *Pushdown processes: Games and Model Checking*, CAV '96, LNCS, vol. 1102, Springer–Verlag, 1996, pp. 62–74.

[Wal00] ———, *Model checking ctl properties of pushdown systems*, FSTTCS '00, LNCS, vol. 1974, Springer–Verlag, 2000, pp. 127–138.

[Wil01] T. Wilke, *Alternating tree automata, parity games, and modal μ–calculus*, Bull. Soc. Math. Belg. **8** (2001), no. 2.

This work is licensed under the Creative Commons Attribution-NoDerivs License. To view a copy of this license, visit http://creativecommons.org/licenses/by-nd/2.0/ or send a letter to Creative Commons, 559 Nathan Abbott Way, Stanford, California 94305, USA.